arXiv:1411.5166v4 [cs.PL] 29 Aug 2016

Subtyping in Java with Generics and Wildcards is a Fractal

Moez A. AbdelGawad
College of Mathematics, Hunan University
Changsha 410082, Hunan, P.R. China
Informatics Research Institute, SRTA-City
New Borg ElArab, Alexandria, Egypt
moez@cs.rice.edu

August 30, 2016

Abstract

While developing their software, professional object-oriented (OO) software developers keep in their minds an image of the subtyping relation between types in their software. The goal of this paper is to present an observation about the graph of the subtyping relation in Java, namely the observation that, after the addition of generics—and of wildcards, in particular—to Java, the graph of the subtyping relation is no longer a simple directed-acyclic graph (DAG), as in pre-generics Java, but is rather a *fractal*. Further, this observation equally applies to other mainstream nominally-typed OO languages (such as C#, C++ and Scala) where generics and wildcards (or some other form of 'variance annotations') are standard features. Accordingly, the shape of the subtyping relation in these OO languages is more complex than a tree or a simple DAG, and indeed is also a fractal. Given the popularity of fractals, the fractal observation may help OO software developers keep a useful and intuitive mental image of their software's subtyping relation, even if it is a little more frightening, and more amazing one than before. With proper support from IDEs, the fractal observation can help OO developers in resolving type errors they may find in their code in lesser time, and with more confidence.

1 Introduction

For helping themselves in writing, debugging and maintaining their software, professional software developers using object-oriented programming languages keep in their minds an image or picture of the subtyping relation between types in their software while they are developing their software. In pre-generics

Java [21], the number of possible object types (also called *reference types*) for a fixed set of classes in a program was a *finite* number (however large it was), and, more importantly, the structure of the subtyping relation between these types (and hence of the mental image a developer kept in mind) was simple: the graph of the subtyping relation between classes and interfaces (*i.e.*, with multiple-inheritance of interfaces) was a simple directed-acyclic graph (DAG), and the graph of the subtyping relation between classes alone (*i.e.*, with single-inheritance only, more accurately called the *subclassing* relation) was simply a tree.[1] This fact about the graph of the subtyping relation applies not only to Java but, more generally, also to the non-generic sublanguage of other mainstream nominally-typed OO languages similar to Java, such as C# [1], C++ [2], and Scala [26].

Today, generics and wildcards (or some other form of 'variance annotations') are a standard feature of mainstream nominally-typed OO languages. The inheritance relation, between classes (and interfaces and traits, in OO languages that support these notions) is still a finite relation, and its shape is still the same as before: a simple DAG. But, given the possibility of arbitrary nesting of generic types, the number of possible object types in a generic Java program has become infinite, and the shape of the subtyping relation in nominally-typed OO languages has become more complex than a tree or a simple DAG.

It is thus natural to wonder, "*What is the shape of the subtyping relation in Java*, now after the addition of generics and wildcards?"

This question on subtyping in Java is similar to one Benoit Mandelbrot, in the 1960s, wondered about: "How long is the coast of Britain?". At that time, some mathematicians (including many computer scientists) used to believe that mathematics was perfect because it had completely banished pictures, even from elementary textbooks. Mandelbrot, using computers, put the pictures back in mathematics, by discovering fractals, and, in the process, finding that Britain's coast has infinite length.

The goal of this paper is to present and defend, even if incompletely and unconventionally (using mainly hierarchy diagrams, and only using equations suggestively), a fundamental observation about the graph of the subtyping relation in Java. We observed that, after the addition of generics—and of wildcards, in particular—to Java, the graph of the subtyping relation is still a DAG, but is no longer a simple DAG but is rather one whose structure can be better understood, of all possibilities, also as a *fractal* – and in fact, as we explain below, an intricately constructed fractal (albeit a different kind of fractal than that of Britain's coast).

To motivate our observation, we use very simple generic class declarations to present in the paper some diagrams for the subtyping relation that represent the iterative construction of the subtyping graph, in the hope of making the construction process very simple to understand and thus make the fractal

[1] Trees are well-known data structures, and a DAG, in short, is a generalization of a tree where a node is further allowed to have more than one parent node (*i.e.*, not just one parent as in a tree) but the node cannot be a parent of itself, even if indirectly; a DAG can thus have no cycles, hence being 'acyclic'.

observation very clear. To further argue for and strengthen the observation, we also suggest algebraic equations for mathematically describing the subtyping fractal and its construction process. (Our equations are akin of recursive domain equations of domain theory that are used to construct 'reflexive domains'. The similarity is suggestive of a strong relationship, possibly even suggesting reflexive domains—useful for giving mathematical meaning for programming languages—might be fractals too, even though we refrain from arguing for this claim here.)

1.1 Practical Significance

Given the popularity fractals enjoy nowadays, we believe the fractal observation about subtyping in nominally-typed OO languages may help OO software developers keep a useful and intuitive mental image of their software's subtyping relation, even if it is a little more frightening, and more amazing one than the one they had before. As an immediate application of the fractal observation, IDEs (Integrated Development Environments) that OO developers use can make developers' lives easier, making them develop their software faster and with more confidence, by presenting to them parts of the fractal representing the subtyping relation in their software and allowing developers to "zoom-in"/"zoom-out" on sections of the fractal/relation that are of interest to the developers, in order for them to better understand the typing relations in their software and so that they may resolve any type errors in their code more quickly and more confidently. OO language designers may also benefit from the fractal observation, since having a better understanding of the subtyping relation may enable them to have a better understanding of the interactions between different features of OO languages—such as the three-tiered interaction, in Java, between generics (including wildcard types), 'lambdas' (formerly known as 'closures') and type inference—leading designers to improve the design of the language, and to better design and implement its compilers. Finally, in allusion to Joshua Bloch's well-known quote when considering adding closures to Java, we hope, by making the fractal observation about subtyping, to enable decreasing (or at least, more accurately estimating) the "*complexity budget*" paid for adding generics and wildcards to Java.

2 The Fractal Observation

As any standard definition (or an image) of a fractal will reveal, fractals (sometimes also called *recursive graphs*, or *self-referential graphs*) are drawings or graphs that are characterized by having "minicopies" of themselves inside of them [25, 14, 3, 4]. Given their *self-similar* nature, when zooming in on a fractal it is not a surprise to find a copy of the original fractal spring up. More generally, the minicopy is not an exact copy, but some *transformation* of the original: it may be the original rotated, translated, reflected, and so on. As such, when constructing a fractal iteratively (as is standard) it is also not a

surprise to add details to the construction of the fractal by using (transformed) copies of the fractal as constructed so far (*i.e.*, as it exists in the current iteration of the construction) to get a better, more accurate approximation of the final fractal (See Figure 2.1, and [4]).

While it may not be immediately obvious to the unsuspecting, but "having transformed minicopies of itself" is exactly what we have noticed also happens in (the graph of) the subtyping relation of Java—and of other similar generic nominally-typed OO languages such as C#, C++, and Scala—after generics and wildcards were added to the Java type system. Figure 2.2 on the following page presents a drawing of the first steps in the construction of a subtyping graph, to illustrate and give a "flavor" of the observation. In Section 3, to motivate presenting the subsequent *transformations observation* in Section 4, we present a more precise and more detailed diagram—one that, unlike Figure 2.2, uses no 'raw types', and has an additional class `D`.

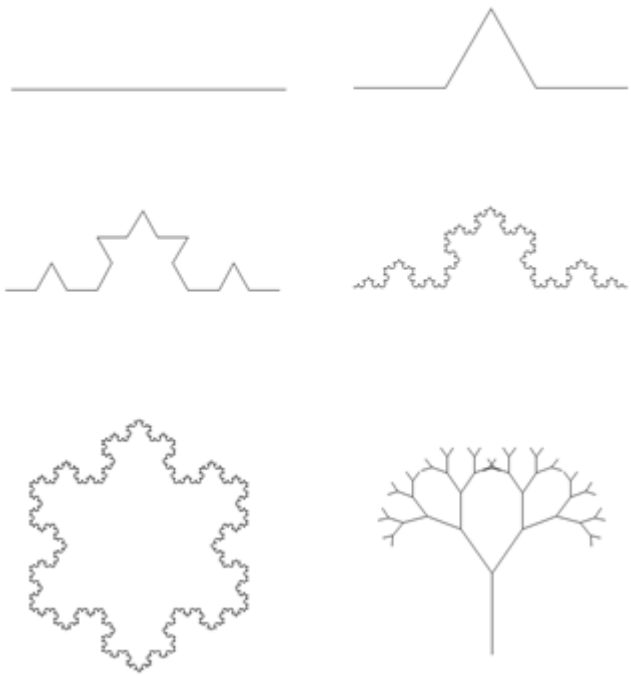

Figure 2.1: Fractals: (First Steps in Constructing) The Koch Curve and (a Step in Constructing) a Fractal Tree

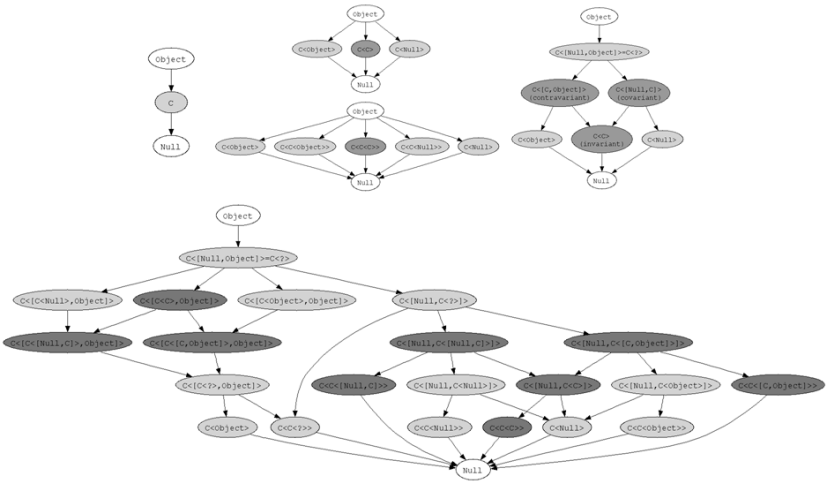

Figure 2.2: First Iterations of Constructing a Subtyping Graph in Java

3 Observation Illustrated: Constructing The Subtyping Fractal

3.1 Subclassing

To illustrate our main observation and how the subtyping fractal is constructed, let us assume we have the non-generic class `Object` (which extends/subclasses no other classes, *i.e.*, is at the top of the subclassing/inheritance hierarchy), and that we have, as expressed in the two simple lines of code below, two generic classes `C` and `D` that extend class `Object` and that take one (unbounded) type parameter. Similarly, and crucial to seeing the subtyping graph as a fractal, we also assume we have a "hidden" (*i.e.*, inexpressible in some OO languages, such as Java) non-generic class `Null` at the bottom of the class inheritance hierarchy (whose only instance is the `null` object, which in Java is an instance of every class and can be assigned to a variable of any object type[2].

```
class C<T> extends Object {}
class D<T> extends Object {}
```

Figure 3.1a on the next page demonstrates the subclassing hierarchy (*a.k.a.*, inheritance hierarchy) based on assuming these class declarations.

[2]Even when the `instanceof` operator in Java, only for developers' convenience, returns `false` as the value of the expression '`null instanceof C`', for any class `C`. Incidentally, the possibility of having *non-nullable* classes, existing in some OO languages such as C#, may need some special provision to envision the resulting subtyping graph, such as providing an additional class `Empty` that extends class `Null` and has no instances. For domain-theorists: The 'class type' corresponding then to class `Empty` will be the empty object domain, *i.e.*, the domain whose only "instance", or member, is $\perp_O$, "the non-terminating object".)

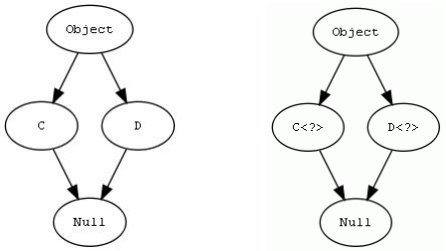

(a) Subclassing (b) Subtyping (Nest. Level 0)

3.2 From Subclassing to Subtyping

3.2.1 Nesting Level 0

The declared inheritance relation between class (and interface/trait) names in a program is the *starting point*[3] for constructing the graph of the subtyping relation in nominally-typed OO languages, including Java (note the use of the identification of type inheritance and subtyping in nominally-typed OOP [5, 6, 8, 7, 18, 12, 13] to interpret 'class extension' as 'subtyping between corresponding class types'. We discuss the role of nominality in more detail in Section 5). Figure 3.1b shows that the "default type argument", namely `?` (the unbounded wildcard type), is used in this initial step as the type argument for all generic classes to form type names for corresponding class types.

3.2.2 Nesting Level 1 ... and Beyond

Figure 3.1 on page 9 demonstrates how the (names of) types in the next iteration of constructing the subtyping fractal (*i.e.*, of the iteration numbered $i+1$, which we can "see" after looking at iteration i if we "zoom in" one step) constructing are constructed by replacing/substituting all the `?`'s in level/iteration 0 (the base step) with *three* different forms of each type `T` in the previous level (level i), namely `? extends T` (covariance), `? super T` (contravariance), and `T` (invariance).[4]

Covariant, contravariant and invariant subtyping rules are then used to decide the subtyping relation between all the newly constructed types (note that,

[3] Or 'the initial graph' or 'the base step' of the recursion "crank," or the 'skeleton' [22] of the graph, or the 'seed' [15] of the fractal.

[4] Which corresponds to defining the fractal using the equation $G = G_0(I(G))$ (see below). Replacing each of the *innermost* (or, all?) `?`s of a type ("holes" in the type) in level i with a `#` (a hash, as a placeholder), then replacing these `#`s with three different forms of each one of the types in the previous level (level i), or in level 0, to construct names of the types of the new level (corresponding to equation $G = G(I(G))$ or $G = G(I(G_0))$). See further comments below for a note on the likely equality of the first two equations, and on the likely uselessness of the third equation as defining the fractal.

due to the inclusion of types `Object` and `Null` in level 0 and in all subsequent levels, all level i types are *also* types of level/iteration $i+1$. This motivates the notion of the rank of a type. The level/iteration in which a type *first* appears is called the *rank* of the type. As such, types `Object` and `Null` are always of rank 0). In Figure 3.1 we use `?xT` and `?sT` as short-hands for `? extends T` and `? super T` respectively.

The effect of variant subtyping rules on the subtyping graph:

- Covariant Subtyping: The level 0 graph is *copied* inside `C<?>` and `D<?>` (illustrated by **green** arrows in diagrams). For ten types `T` (from 2 non-generic classes + 2 generic classes × 4 types in level 0), we have paths
  `Object -> C<? extends T> -> Null`, and
  `Object -> D<? extends T> -> Null`
  (Note: `? extends Null` is the same as `Null`. Inexpressible in Java).

- Contravariant Subtyping: The level 0 graph is *flipped* (turned upside-down) inside `C<?>` and `D<?>` (illustrated by **red** arrows in diagrams). For ten types `T`, like for covariance, we have paths
  `Object -> C<? super T> -> Null`, and
  `Object -> D<? super T> -> Null`
  (Note: `? super Object` is the same as `Object`. See footnote 5 on the following page regarding current Java behavior).

- Invariant Subtyping: The level 0 graph is *flattened* inside `C<?>` and `D<?>` (no corresponding arrows in diagrams). For ten types `T`, like for covariance, we have paths
  `Object -> C<T> -> Null`, and
  `Object -> D<T> -> Null` .

3.3 Type Intervals

Figure 3.2 on page 9 illustrates how to use the notion of type intervals [11, 10] to combine all three (*i.e.*, covariant, contravariant and invariant) subtyping rules (and to add even more types to the subtyping relation in later iterations/nesting levels).

In Figure 3.2, we have all three transformations applied to level 0 graph and embedded inside `C<?>` and `D<?>` (note that bounds of an interval can degenerately be equal types, corresponding to invariance). For twenty types `S` and `T` (where `S` is a subtype of `T` in the previous iteration/level), from 2 non-generic classes + 2 generic classes × 9 intervals in level 0, we have `Object -> C<S-T> -> Null`, `Object -> D<S-T> -> Null` (The notation `S-T` means the interval with lowerbound `S` and upperbound `T`. For brevity, we use `O` for `Object` and `N` for `Null`). If class `C` or class `D` had subclasses other than `Null`, this graph diagram would have been even richer—*i.e.*, it would have had more types—than

the graph in Figure 3.1 on page 9. (It can be noted that the `Null` type is useful in expressing intervals. Yet the diagram can be presented without it, using `extends` only or `super` only, while allowing but not requiring a naked `?`; or, for brevity, using a symbol like `<:`).

Note: The types `Null`, `C<Null>`, and `D<Null>`, inside dotted graph nodes in Figure 3.1 on the next page and Figure 3.2 on the following page, are currently *inexpressible* in Java (*i.e.*, as of Java 8, based on the assumption that these types are of little practical use to developers.) Subtyping relations involving these inexpressible types are also currently of little use to Java developers (except in type inference). Accordingly, they also are drawn in Figure 3.2 using dotted graph edges.[5]

It should now be clear how to constructing the rest of the subtyping fractal. Each next nesting level of generics corresponds to "zooming one level in" in the subtyping fractal, and the construction of the new "zoomed-in" graph is done using the same method above, where wildcards (or, intervals) over the previous subtyping graph substitute all the `?` in that graph to produce the next level graph of the subtyping relation. And there is nothing in generics that disallows arbitrarily-deep, potentially infinite, nesting.

4 The Transformations Observation

While making the fractal observation, we made yet another observation that helps explain the fractal observation more deeply. In particular, we noted that in constructing the graph of the subtyping relation, when moving from types of a specific level of nesting to types of the next deeper level (*i.e.*, when "zooming in" inside the graph of the relation, or when doing the inductive step of the recursive definition of the graph), *three* kinds of *transformations* are applied to the level i subtyping graph, in agreement with the general nature of fractals having transformed minicopies of themselves embedded within. We call these three transformations the *identity* (or, *copying*) transformation, the *upside-down* reflection (or, *flipping* transformation), and the *flattening* transformation. The first transformation (identity) makes an exact copy of the input subtyping relation, the second transformation (upside-down reflection) flips over the relation (a subtype in the input relation becomes a supertype, and vice versa), while the third transformation "attempts to do both (*i.e.*, the identity and flipover

[5] (**A bug in Javac**): Also, as of Java 8, we have noted that Java does not currently identify `? super Object` with `Object`, and as such a variable `b` of type `C<? super Object>`, for example, *cannot* be assigned to a variable `a` of type `C<Object>` (*i.e.*, for the statement `a=b;` the Java compiler `javac` currently emits a type error with an unhelpful semi-cryptic error message that involves 'wildcard capturing') even as Java allows the opposite assignment of `a` to `b` (*i.e.*, the statement `b=a;`), implying that, even though Java currently correctly sees `C<Object>` as a subtype of `C<? super Object>`, it currently does *not* consider `C<? super Object>` as a subtype of `C<Object>`. Given that there are no supertypes of type `Object` (the class type corresponding to class `Object`), and it is not expected there will ever be any, we believe the Java type system should be fixed to identify the two type arguments `? super Object` and `Object`, and thus correctly allow the mentioned currently-disallowed assignments.

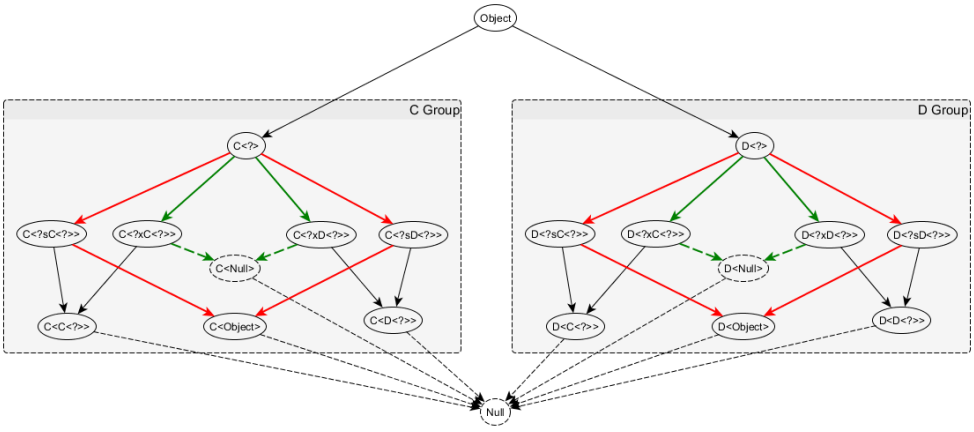

Figure 3.1: Subtyping: (First) Inductive Step. Nesting Level 1

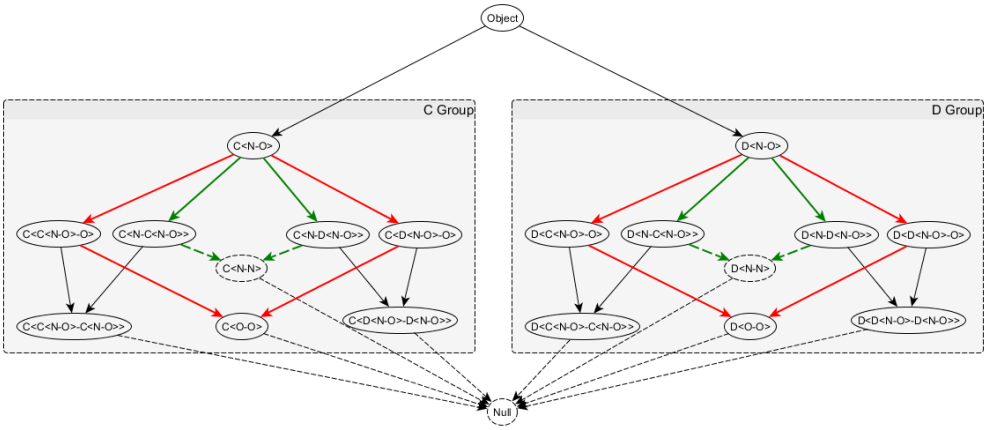

Figure 3.2: Subtyping Generalization with Intervals. Nesting Level 1

transformations)," in effect making types that were related in its input subtyping relation be *unrelated* in its output subtyping relation (hence the output of this transformation is a "flat" relation, called an *anti-chain*.)

Explaining this observation regarding the subtyping fractal in terms of OO subtyping is done by noting that the three mentioned transformations correspond to (in fact, result from) the covariant subtyping rule, contravariant subtyping rule, and invariant subtyping rule, respectively. This is demonstrated, in a very abridged manner, in Figure 3.1 on the previous page (with the green arrows corresponding to copying the previous level graph, corresponding to covariant subtyping, the red arrows corresponding to flipping over the previous level graph, corresponding to contravariant subtyping.) It should be noted that *also* the level 1 graph as a *whole* is the same structure as the level 0 graph when the 'C Group' nodes are lumped into one node and the same for the 'D Group' node. That means that, in agreement with the graph being a fractal (where self-similarity must exist at all levels of *scale*), when the graph of subtyping is "viewed from far" it *looks the same* as the level 0 graph. In fact, when looked at from a far enough distance this similarity to the level 0 graph will be the case for all level i , where $i \geq 1$, graphs.

5 Nominally-typed OOP vs. Structurally-typed OOP

It should be noted that class names information (*a.k.a.*, nominality, and 'nominal type information') of nominally-typed OO languages (such as Java, C#, C++, and Scala) is used in the base/first construction step in constructing the subtyping relation between generic types as a fractal. In contrast, structurally-typed OO languages (such as OCaml [23], Moby [20], PolyTOIL [17], and StrongTalk [16]), known mainly among programming languages researchers, do *not* have such a simple base step, since a record type corresponding to a class (with at least one method) in these languages does *not* have a finite number of supertypes to begin with, given that "superclasses of a class" in the program, when viewed structurally as supertypes of record types, do *not* form a finite set. Any record type has an infinite set of record subtypes (due to their width-subtyping rule [28]). Accordingly, a record type with a method—*i.e.*, a member having a function type—causes the record type to have an infinite set of *supertypes*, due to contravariance of the type of the method. Adding-in a depth-subtyping rule makes the subtyping relation between record types with functional member types even more complex.

This motivates suspecting that subtyping in structurally-typed OO language is a *dense* relation, in which every pair of non-equal types in the relation has a third type, not equal to either member of the pair, that is "in the middle" between the two elements of the pair, *i.e.*, that is a subtype of the supertype (the upperbound) of the pair and a supertype of the subtype (the lowerbound) of the pair. In fact this may turn out to be simple to prove. Due to a class

in generic nominally-typed OO languages having a finite set of superclasses in the subclassing relation, subtyping in generic nominally-typed OO languages languages is not an (everywhere) dense relation, and the subclassing relation in these languages forms a simple finite basis (the "skeleton") for constructing the subtyping relation. For structurally-typed OO languages (where record types with functional members are a must, to model structural objects), this basis (the "skeleton") is infinite and thus the "fractal" structural subtyping graph (if indeed it is a fractal) is not easy to draw or to even imagine.

For more details on the technical and mathematical differences between nominally-typed and structurally-typed OOP, the interested reader may consult [19, 24, 27, 5, 6, 7, 8, 18, 13, 12, 9, 11].

6 Conclusion

In this paper we presented an observation connecting subtyping in generic nominally-typed OO languages to fractals. We presented diagram for graphs of the subtyping relation demonstrating the iterative process of constructing the relation as a fractal. We further made an observation connecting the three variant subtyping rules in generic OOP to the three transformations done on the graph of the relation for embedding inside the relation. We further noted some possible differences between generic nominally-typed OOP and polymorphic structurally-typed OOP as to the fractal nature of their subtyping relations. (See the Appendix for some further notes, observations and conclusions that may be built on top of the observations and discussions we made, including a suggestive discussion on the use of algebraic equations to precisely describe the generic OO subtyping relation as a fractal).

References

[1] C# language specification, version 3.0. http://msdn.microsoft.com/vcsharp, 2007.

[2] *ISO/IEC 14882:2011: Programming Languages: C++*. 2011.

[3] Nova | hunting the hidden dimension - pbs. http://www.pbs.org/wgbh/nova/physics/hunting-hidden-dimension.html, 2011.

[4] Fractal. http://en.wikipedia.org/wiki/Fractal, Dec. 2014.

[5] Moez A. AbdelGawad. *NOOP: A Mathematical Model of Object-Oriented Programming*. PhD thesis, Rice University, 2012.

[6] Moez A. AbdelGawad. *NOOP: A Nominal Mathematical Model Of Object-Oriented Programming*. Scholar's Press, 2013.

[7] Moez A. AbdelGawad. An overview of nominal-typing versus structural-typing in object-oriented programming (with code examples). Technical report, arXiv.org:1309.2348 [cs.PL], 2013.

[8] Moez A. AbdelGawad. A domain-theoretic model of nominally-typed object-oriented programming. *Journal of Electronic Notes in Theoretical Computer Science (ENTCS), DOI: 10.1016/j.entcs.2014.01.002.*, 301:3–19, 2014.

[9] Moez A. AbdelGawad. A comparison of NOOP to structural domain-theoretic models of object-oriented programming. *Preprint available at http://arXiv.org/abs/1603.08648*, 2016.

[10] Moez A. AbdelGawad. Towards an accurate mathematical model of generic nominally-typed OOP (extended abstract). 2016.

[11] Moez A. AbdelGawad. Towards understanding generics. Technical report, arXiv:1605.01480 [cs.PL], 2016.

[12] Moez A. AbdelGawad. Why nominal-typing matters in OOP. *Preprint available at http://arxiv.org/abs/1606.03809*, 2016.

[13] Moez A. AbdelGawad and Robert Cartwright. In nominally-typed OOP, objects are not mere records and inheritance *Is* subtyping. *Submitted for journal publication*, 2016.

[14] Michael F Barnsley. *Fractals Everywhere: New Edition*. Courier Dover Publications, 2013.

[15] Richard Bird et al. *Introduction to functional programming using Haskell*, volume 2. Prentice Hall Europe Hemel Hempstead, UK, 1998.

[16] G. Bracha and D. Griswold. Strongtalk: typechecking Smalltalk in a production environment. In *OOPSLA'93*, pages 215–230, 1993.

[17] K. Bruce, A. Schuett, R. van Gent, and A. Fiech. PolyTOIL: A type-safe polymorphic object-oriented language. *ACM Transactions on Programming Languages and Systems*, 25(2):225–290, 2003.

[18] Robert Cartwright and Moez A. AbdelGawad. Inheritance *Is* subtyping (extended abstract). In *The 25th Nordic Workshop on Programming Theory (NWPT)*, Tallinn, Estonia, 2013.

[19] Robert Bruce Findler, Matthew Flatt, and Matthias Felleisen. Semantic casts: Contracts and structural subtyping in a nominal world. In *ECOOP 2004–Object-Oriented Programming*, pages 365–389. Springer, 2004.

[20] K. Fisher and J. Reppy. The design of a class mechanism for Moby. In *PLDI*, 1999.

[21] James Gosling, Bill Joy, Guy Steele, and Gilad Bracha. *The Java Language Specification*. Addison-Wesley, 2005.

[22] Douglas R. Hofstadter. *Gödel, Escher, Bach: an Eternal Golden Braid*. Basic Books, second edition, 1999.

[23] X. Leroy, D. Doligez, J. Garrigue, D. Rémy, and J. Vouillon. The Objective Caml system. Available at http://caml.inria.fr/.

[24] Donna Malayeri and Jonathan Aldrich. Integrating nominal and structural subtyping. In *ECOOP 2008–Object-Oriented Programming*, pages 260–284. Springer, 2008.

[25] Benoit B Mandelbrot. Fractals. 1977.

[26] Martin Odersky. The scala language specification, v. 2.7. http://www.scala-lang.org, 2009.

[27] Klaus Ostermann. Nominal and structural subtyping in component-based programming. *Journal of Object Technology*, 7(1):121–145, 2008.

[28] Benjamin C. Pierce. *Types and Programming Languages*. MIT Press, 2002.

[29] David I Spivak. *Category theory for the sciences*. MIT Press, 2014.

A Further Notes

The following notes and observations can be added to the ones we made in the main paper:

1. Relations on Type Intervals: For type intervals $I = [S, T]$, where $S <: T$, with lowerbound (S) and upperbound (T), two relations on intervals can be defined that can help in constructing the subtyping fractal: An interval containing another interval (the *contains* relation: $S_1 <: S_2 \wedge T_2 <: T_1$), and an interval preceding another interval (the *precedes* relation: $T_1 <: S_2$).

2. The *Pruning Transformation*: Bounds, *i.e.*, lowerbounds or upperbounds, on a type parameter limit (*i.e.*, decrease) the types of level i that can substitute the holes (the `?`s) when constructing a type in level $i + 1$, so pruning means that a substitution *respects* these declared bounds.

3. Demonstration Software: An interactive Mathematica program that demonstrates the iterative construction of the subtyping hierarchy, for multiple simple class hierarchies, up to four nesting levels is available upon request (The program uses the `Manipulate` function of Mathematica 6, is formatted as a Mathematica 6 demo, and is in the Mathematica .nb format, *i.e.*, the file format Mathematica has used as of 2007.)

4. Multi-arity: Generic classes with *multiple* type parameters simply result in types with multiple "holes" at the same nesting level for the same class.

5. Graph Matrices: Representing successive subtyping graphs as adjacency matrices (0-1 matrices) is useful in computing (paths in graph of) the relation (and in computing containment of intervals). (Using $(I - A)^{-1}$, with binary addition and multiplication of matrices, to compute the transitive closure of the relation and thus paths/intervals over it).

6. Category Theory: Given the use of the notion of *operads* in category theory to model self-similarity [29], we intend to consider the possibility of using operads to express and communicate the fractal nature of the generic OO subtyping relation.

7. Algebraic Equations: According to Benoit Mandelbrot, Hermann Well wrote that 'the angel of geometry and the devil of algebra share the stage, illustrating the difficulties of both.' Turning to some algebra, we expect the graph of the subtyping relation to be described by a recursive equation, as is the case for many fractals. We anticipate this equation to be (something along the lines of)

$$G = G_0(copy(G) + flip(G) + flatten(G)),$$

where G_0 stands for the initial graph (the 'skeleton' of the subtyping fractal, resulting from turning the subclassing relation into a subtyping relation by using `?` as the default type argument for generic classes), and the application of G_0 to its argument (another graph) means the *substitution* (similar to β-reduction in λ-calculus) of its "holes" (the `?`'s in its types/nodes) with the argument graph (*i.e.*, the graph $copy(G) + flip(G) + flatten(G)$ which applies the three above-mentioned transformations to G, and where $+$ means "subtyping-respecting union" of component graphs.)

1. More on Algebraic Equations:

   (a) The G_0 in the equation above—*i.e.*, the graph of the first iteration of the subtying relation, which is directly based on the subclassing relation—is what makes (all iterations/approximations of) the graph G have the same structure "when viewed from far", *i.e.*, when zooming out of it, as the subclassing relation).

   (b) To construct approximations of G iteratively, the equation can be interpreted to mean

$$G_{i+1} = G_0(copy(G_i) + flip(G_i) + flatten(G_i)),$$

   which means when constructing approximations to G we construct elements of the sequence

$$G_0 = G_0,$$

$$G_1 = G_0(copy(G_0) + flip(G_0) + flatten(G_0)),$$

$$\begin{aligned} G_2 &= G_0(copy(G_1) + flip(G_1) + flatten(G_1)) \\ &= G_0(copy(G_0(copy(G_0) + flip(G_0) + flatten(G_0))) + \\ &\quad flip(G_0(copy(G_0) + flip(G_0) + flatten(G_0))) + \\ &\quad flatten(G_0(copy(G_0) + flip(G_0) + flatten(G_0)))) \end{aligned}$$

$$G_3 = G_0(copy(G_2) + flip(G_2) + flatten(G_2)) = ...,$$

... etc.

(c) Another seemingly-equivalent recursive equation for describing the subtyping graph G is

$$G = G(copy(G_0) + flip(G_0) + flatten(G_0)),$$

which, even though not in the more familiar $x = f(x)$ format, has the advantage of showing that G (the limit, infinite graph) is equivalent to (isomorphic to) substituting its own holes with transformations of G_0, *i.e.*, that the substitution does *not* affect the final infinite graph G (just as adding 1 to ω, the limit of natural numbers, does not affect its cardinality; $|\omega| = |\omega + 1|$.) It also reflects the zooming-in fact (opposite to the zooming-out fact above) that when zooming-in into G we find (transformed copies of) G_0 each time we zoom in, ad infinitum. (See Note 5 below for why we believe this third equation may in fact be *incorrect*.)

2. Algebraic Equations with Intervals: With intervals, the equation above becomes simpler and more general, where, if I is the function computing all the intervals over a graph, we then have

$$G = G_0(I(G)), \text{ or, } G = G(I(G_0)),$$

or, most accurately,

$$G = G(I(G)).$$

Note that the three equations agree on defining $G_1 = G_0(I(G_0))$. The three equations disagree however on later terms of the construction sequence. They, for example, define $G_2 = G_0(I(G_1))$, $G_2 = G_1(I(G_0))$, and $G_2 = G_1(I(G_1))$, respectively. The equivalence of the three equations (*i.e.*, of the resulting graph from each) is unlikely, but a mathematical proof or a convincing intuitive proof of that is needed (see Note 5 below, however).

3. Benefits and Applications: An obvious benefit of the observation in this paper is to demonstrate one more (unexpected?) place where fractals show

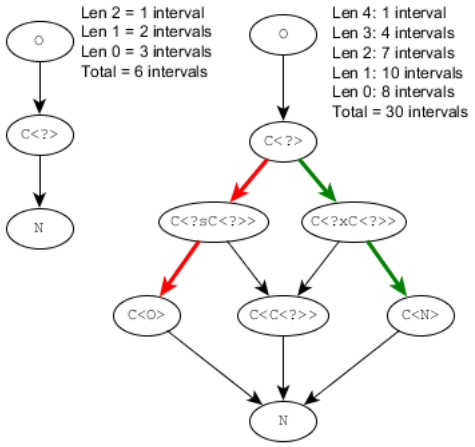

Figure A.1: Subtyping levels 0, 1 with one generic class (`C`)

up. Yet an additional benefit, and practical application, of the observation may be to apply some of the theory developed for fractals to better the understanding of the subtyping relation in OO languages, possibly leading to providing a better understanding of their generic type systems and thus developing better OO language compilers.

4. Parameterizing classes `Object` and `Null`: At least one needs to be non-parameterized, if not both? Otherwise we may have an unbounded infinite ascending chain of supertypes (see Section 5.) (What will then be the meaning of `?`, and be the default type argument?)

5. More Levels/Iterations: To further demonstrate the fractal observation, and to help resolve which of the three equations above (best) describes the graph of the subtyping relation, we draw the level 2 graph G_2 using a simpler initial graph (*i.e.*, the 'skeleton') than we used for the earlier figures. See Figure A.1 and Figure A.2 on the next page.

   Some notes on G_0, G_1, G_2:

   (a) G_2 , constructed as $G_0(I(G_1))$, has 32 nodes, and 66 edges.

   (b) The number of levels in graphs G_0, G_1, G_2, ... (*i.e.*, the maximum path length) increases by two each time (2, 4, 6, 8, ...). This is clear in the diagrams, particularly ones with colored arrows.

   (c) The number of nodes and edges in G_0, G_1, G_2, ...:

      i. Nodes: G_0**=3**=(2+1), G_1**=8**=(2+3+2+1), G_2**=32**=(2+8+10+7+4+1), ... **???**=(2+32+66+<...4 numbers...>+1).

      ii. Edges: G_0**=2**, G_1**=10**, G_2**=66**, ... **???**.

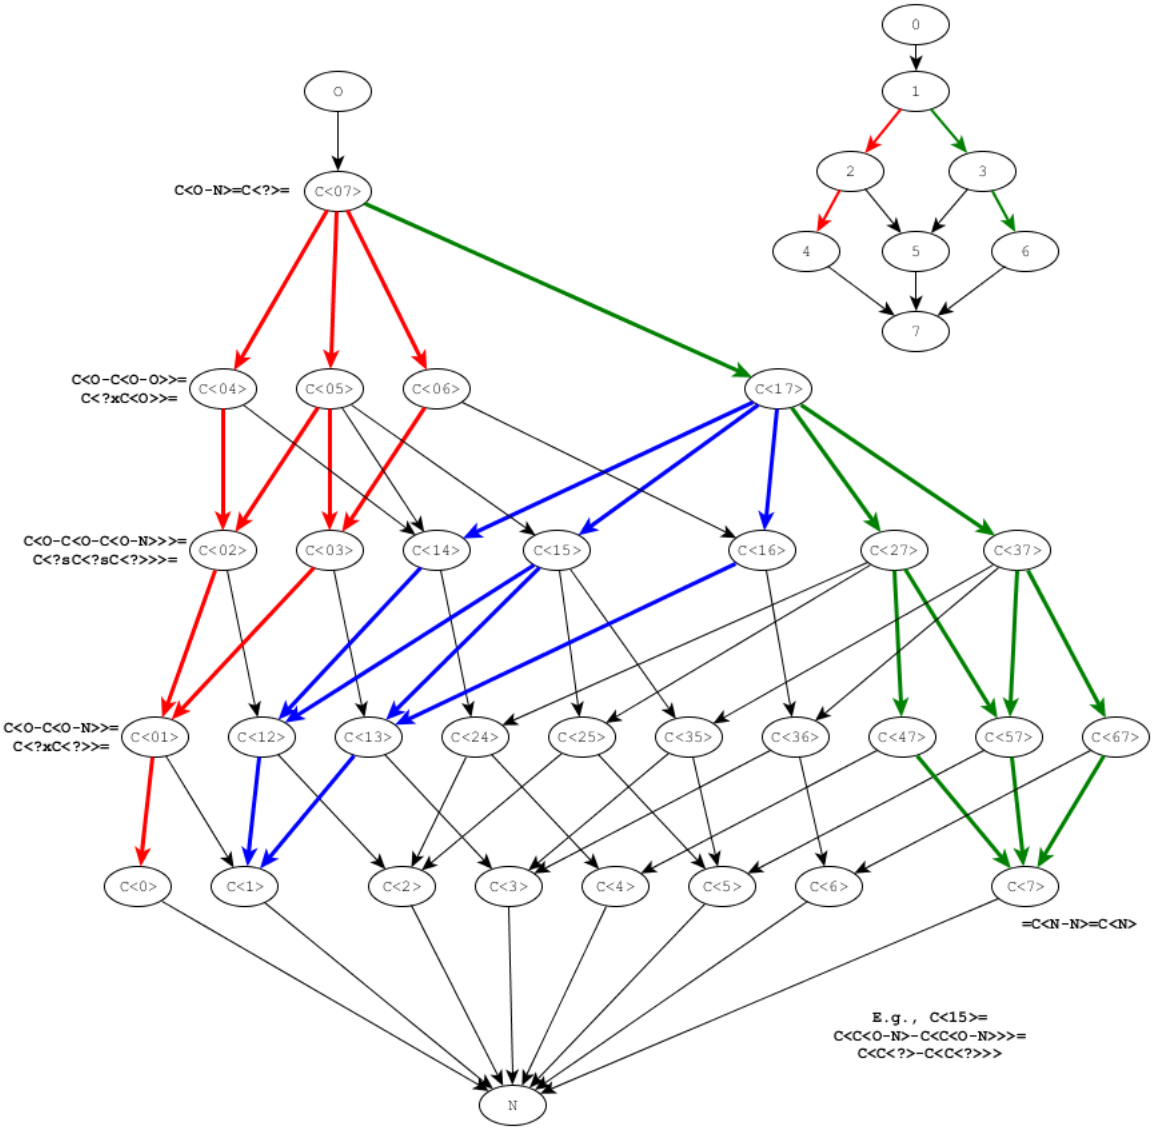

Figure A.2: Subtyping levels 1, 2, corresponding to equation $G = G_0(I(G))$ and Bounded Contravariance/Graph-Flip (Highlighted in Blue)

(d) Algebraic Equations:

i. G_2, as constructed above, is the same graph as $G_1(I(G_1))$ ... !!

ii. Thus, $G_0(I(G_1)) = G_1(I(G_1))$, meaning that, given $G_1 = G_0(I(G_0))$, we have

$$G_0(I(G_0(I(G_0)))) = G_0(I(G_0))\,(I(G_0(I(G_0))))$$

iii. The skeptic reader may trying constructing the graph corresponding to the equation $G_2 = G_1(I(G_1))$

iv. **Proof**: Each type/node constructed in $G_1(I(G_1))$ is constructed in $G_0(I(G_1))$ (and vice versa, which is easy to see).

A. Same for proving $G_1(I(G_0)) = G_0(I(G_0))$, which means we have

$$G_0(I(G_0))\,(I(G_0)) = G_0(I(G_0))$$

v. (An Analogy) Something unknown becoming known. Knowing it again does not add new.

vi. Philosophical observation, using 'Old' = G_i, 'New' = G_{i+1}:

A. New in New = New in Old

B. Old in New != New in Old

C. Old in Old = New = Old in New

(e) In addition to subgraphs highlighted in green and red (which show an exact copy and a flipped copy, due to covariance and contravariance respectively) of G_1 inside G_2, Figure A.2 on the preceding page also shows a miniature *pruned* flipped copy of G_1 inside G_2, highlighted in blue (due to bounded contravariance).